

T E S T I M O N Y

RAND

The Role of Technical Standards in Today's Society and in the Future

Martin C. Libicki

CT-172

September 2000

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

Science and Technology Policy Institute

The RAND testimony series contains the statements
of RAND staff members as prepared for delivery.

RAND is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. RAND® is a registered trademark. RAND's publications do not necessarily reflect the opinions or policies of its research sponsors.

Published 2001 by RAND

1700 Main Street, P.O. Box 2138, Santa Monica, CA 90407-2138

1200 South Hayes Street, Arlington, Virginia 22202-5050

RAND URL: <http://www.rand.org/>

To order RAND documents or to obtain additional information, contact RAND Distribution Services (Telephone: 310-451-7002; Fax: 310-451-6915; or Internet: order@rand.org). Abstracts of all RAND documents may be viewed on the World Wide Web (<http://www.rand.org>). Publications are distributed to the trade by National Book Network.

**STATEMENT OF
Martin C. Libicki
Senior Policy Analyst
RAND**

**Before the
Committee on Science
Subcommittee on Technology
United States House of Representatives**

Hearing on the Role of Technical Standards in Today's Society and in the Future

September 13, 2000

RAND is a nonprofit institution that helps improve policy and decisionmaking through research and analysis. This statement is based on a variety of sources, including research conducted at RAND. However, the opinions and conclusions expressed are those of the author and should not be interpreted as representing those of RAND or any of the agencies or others sponsoring its research.

20010308 021

Good morning, Madame Chairwoman, and thank you for inviting me to testify. I am Martin Libicki, of Kensington, Maryland. I work for RAND, a nonprofit institution that helps improve policy and decisionmaking through research and analysis. This statement is based on a variety of sources, including research conducted at RAND. However, the opinions and conclusions expressed are mine and should not be interpreted as representing those of RAND or any of the agencies or others sponsoring its research. My full written testimony consists of chapters from the book I co-authored, **Scaffolding the New Web: Standards and Standards Policy for the Digital Economy**.

Information technology is changing fast. Not surprisingly, standards -- the methods used to bring some order to the process -- have changed apace. Standards can be, yes, boring and even annoying. They are often accused of impeding innovation and originality. But the world of information technology would never have left the blackboards without them. Computing and communications is no more possible without common conventions than is conversation without a common language. But once there is agreement on the modalities of interchange, there are no limits to the kind of creativity that can be expressed using such common terms.

Until roughly 1990, standards development was regular and straightforward. The primary tension was between closed proprietary and open public standards. The latter were developed by voluntary organizations spawned by or at least coordinated with the American National Standards Institute in this country. International imprimatur came from ISO, the International Organization for Standards. Communications standards, in turn, were handled through the International Telecommunications Union, a UN-treaty organization, with representatives from national governments.

Several factors threw this formal system into disarray: the blurring of computation and communication which raised issues that fell in neither camp, the rise of the PC and as a consequence the rise of major players not grown up in the system. But the Open Systems Interconnection (OSI) debacle did not help. It turned out to be very hard to get customers and suppliers to line up behind an ambitious set of formal data communications standards. Meanwhile, off to the side, a number of Government-financed engineers working under the auspices of the Internet Engineering Task Force (or IETF) developed a simpler, but good enough, set of data communications protocols, and, a functioning internetwork, the Internet. By 1993, the tide in standards had shifted; the IETF was the way to go.

Ironically, the IETF acquired its laurels just as its ability to generate new standards slowed down. Or not so ironically. The IETF could work its build-a-little-test-a-little rough-and-ready consensus model only as long as meetings were small, everyone knew each other, and, most importantly, no serious money was involved. Once the Internet was "discovered", corporate interest and meeting size grew. Standards emerged far more slowly. This was not really the fault of the IETF; indeed, it had been anticipated as far back as 1992.

Meanwhile, people like Tim Berners-Lee, and Mark Andreessen were inventing something called the World Wide Web -- a good idea that captured computerdom's imagination. But the Web's protocols (such as URLs, HTML, HTTP, and GIF) became de facto standards without an IETF stamp. Indeed, not until mid-1999 did HTTP become a full IETF standard; a status not then reached by URLs.

Today, the ecology of standards development is less orderly but nevertheless robust and variegated: from treaty organizations such as the ITU, to the formal apparatus under ISO, well-established entities such as the IETF, large continuing consortia such as the World Wide Web Consortium that developed the Extensible Markup Language (XML), essentially one-shot consortia such as those which developed the Wireless Access Protocol (WAP), charismatic networks such as those building Linux, informal teaming arrangements (such as the Open Financial Exchange) for E-commerce, corporations such as Sun (of Java fame) or Adobe (which did Postscript) that hope to raise their inventions into quasi-open standards, and corporations such as Microsoft, that market de facto standards.

Fundamentally, therefore, our research suggests that the standards ecology is healthy and is capable of handling the foreseeable next steps into E-commerce.

Nevertheless, two problems loom. Standardizers are worried about liberally defined software patents. These threaten to make it difficult to use standards without paying licensing fees to someone, often a someone that no one working on the standard had heard from. The other challenge is coming up with standard terms for E-commerce content now that there is general agreement about the grammar -- XML. That is, everyone recognizes how to mark out "price" in a document, but there is no agreement about exactly what "price" means: is it wholesale or retail, dependent on quality-of-service, subject to which terms and limitations, and so on. Hence the cliché: the wonderful thing about standards is how many there are to choose from. E-commerce consortia are constantly being formed and a veritable Babel of tag sets is emerging.

Finally, as for the function of the National Institute of Standards and Technology, it should keep doing what it traditionally does, only more so and better. Three roles suggest themselves: providing an expertly facilitated neutral meeting ground for the development of consensus, developing test methods by which standards and conformance to standards can be measured, and acting as a clearinghouse for standards development, particularly in the E-commerce arena.

Martin C. Libicki
Biography

Dr. Libicki has been a Senior Policy Analyst at RAND (since 1998) working on the relationship between information technology and national security. He has written on Information Technology Standards (*Scaffolding the New Web: Standards and Standards Policy for the Digital Economy* [RAND, 2000], *Information Technology Standards: Quest for the Common Byte* [Digital Press, 1995], the Revolution in Military Affairs (*Illuminating Tomorrow's War* [NDU, 1999], and *The Mesh and the Net* [NDU, 1994]), and on Information Warfare (What is Information Warfare [NDU, 1995] and *Defending Cyberspace and Other Metaphors* [NDU, 1997]). Prior to RAND Dr. Libicki was a senior fellow at NDU's Institute for National Strategic Studies for a dozen years. He received his Ph.D. at Berkeley in 1978 for work on industrial economics.



Document No: MR-1215-OSTP

Title: Scaffolding the New Web: Standards and Standards Policy for the Digital Economy.

Author(s): Martin Libicki, James Schneider, David Frelinger, Anna Slomovic

Abstract: Although much of the growing digital economy rests on the Internet and World Wide Web, which in turn rest on information technology standards, it is unclear how much longer the current momentum can be sustained absent new standards. To discover whether today's standards processes are adequate, where they are taking the industry, and whether government intervention will be required to address systemic failures in their development, RAND undertook five case studies. So far, it seems, the current standards process remains basically healthy, with various consortia taking up the reins of the process, and the rise of open-source software has also aided vendor-neutral standardization. Nevertheless, the prospects for semantic standards to fulfill XML's promise are uncertain. Can the federal government help? Its policy on software patents clearly merits revisiting. More proactively, the National Institute for Standards and Technology could intensify its traditional functions: developing metrologies; broadening the technology base; and constructing, on neutral ground, terrain maps of the various electronic-commerce standards and standards contenders.

SUMMARY

With every passing month, the digital economy grows stronger and more attractive. Much, perhaps, most of this economy rests upon the Internet and its World Wide Web. They, in turn, rest upon information technology standards. Today's standards appear good enough to see the digital economy through the next few years. But it is unclear how much longer the momentum of such commerce can be sustained absent new standards. Are today's standards processes adequate? Where are they taking the industry (and where is the industry taking them)? Is government intervention required to address systemic failures in their development?

To answer these questions, a RAND Science and Technology Policy research team undertook five case studies covering

1. existing Web standards
2. the extensible markup language, XML
3. digital library standards
4. issues related to payments, property, and privacy
5. evolving electronic commerce value chains.

A White House-sponsored meeting of standards practitioners also generated material helpful in developing an overall assessment. All this material was used to inform the body of the report.

Information technology standards are a means by which two or more products (or systems) can function together. Some standards permit peers to interoperate or to exchange data in ways that are mutually

comprehensible. Others permit one thing (e.g., a software application) to work atop another (e.g., an operating system). Information technology has seen a long march away from proprietary conventions (e.g., how the alphabet is translated into bit strings) toward open conventions that have become standards. The Internet and the World Wide Web exemplify openness; their standards are public and largely vendor-neutral. Yet as more products follow standards, innovative products are, almost by definition, unstandardized (in communications, standards often precede product introduction: One phone is useless by itself). So, the conflict between different ways of doing things starts anew. Standards failures tend to have one of three consequences:

1. New activities are stillborn.
2. New activities emerge, but with little interoperability among domains (each with its own conventions).
3. Proprietary standards enable an active but biased marketplace, reducing competition and hobbling innovation.

So far, the process by which standards are written and stamped remains basically healthy. True, the formal standards development organizations that were overtaken by the Internet Engineering Task Force (IETF) in the early 1990s remain on the periphery of the process; the IETF itself has become congested by its own popularity. But consortia (e.g., the World Wide Web Consortium) and forums (e.g., the Wireless Access Protocol Forum) appear to have picked up the slack. The rise of open-source software (e.g., Linux, Apache, Mozilla) has been another force for vendor-neutral standardization.

Yet, the case studies suggest that the success of standards in the marketplace depends on the play of larger forces. HTML and, to a lesser extent, Java succeeded because they were straightforward and unique ways of doing interesting things. But today's Web standards developments are wrapped up in the contests between corporations waging wars over browsers and other Web on-ramps, each trying to do an end-run around each other's proprietary advantages. The standards that would govern digital libraries, intellectual property rights, payments, and privacy are buffeted by the varied interests of affected groups—authors, librarians, rights holders, consumers, banks, merchants, privacy activists, and governments. Although

XML has quickly achieved wide acceptance, it is only the grammar through which Web content can be described. Many groups now vie to establish the words (i.e., the tag sets) everyone else will use: The result so far is a high head of froth and thin beer beneath.

The battle over tag sets reflects the broader problem of describing the messy real world to the sheltered naïfs that our computers still are. There is no obvious way to achieve semantic standardization. Creating one master tag set is optimal but a long shot. Creating tag sets specialized for various communities may be only somewhat more likely but complicates communicating *across* domains (each of which then also needs its own software). Translators would obviate the need for standards, but reliable translation exceeds what today's technology can provide. Ontologies into which everyone's terms can be mapped might improve translation, but how will a standardized ontology come about? Perhaps the best outcome is that some terms are globally standardized; some are locally standardized; and the rest are anyone's guess. There is, incidentally, little cry for the U.S. government to dictate what tags to use.

Does government, in fact, have much of a role to play? Standards for describing and measuring content (e.g., movie ratings, cyber-security performance) may substitute for some regulation. But less may be more: Many standards developers already believe that the government's overly liberal granting of patents on software (and business processes) frustrates the development of standards. Researchers might be allowed to use a fraction of their government research and development funding to work on standards. Perhaps the best help the government can offer is to have the National Institute for Standards and Technology (NIST—specifically, its Information Technology Laboratory) intensify its traditional functions: developing metrologies; broadening the technology base; and constructing, on neutral ground, terrain maps of the various electronic-commerce standards and standards contenders.

Openness is an underlying technical and philosophical tenet of the expansion of electronic commerce. The widespread adoption of the Internet as a platform for business is due to its non-proprietary standards and open nature as well as to the huge industry that has evolved to support it. The economic power that stems from joining a large network will help to ensure that new standards will remain open. More importantly, openness has emerged as a strategy, with many of the most successful e-commerce ventures granting business partners and consumers unparalleled access to their inner workings, databases, and personnel. This has led to a shift in the role of consumers, who are increasingly implicated as partners in product design and creation. An expectation of openness is building on the part of consumers [and] citizens, which will cause transformations, for better (e.g., increased transparency, competition) or for worse (e.g., potential invasion of privacy), in the economy and society.

—Organisation for Economic Cooperation
and Development, 1999

The digital economy sits at the uneasy juncture that separates the idealism of its youth from the moneymaking of its maturity. As a whole, it is *terra incognita*: Everything is new; the landscape is sure to change even as it is brought under the plow; and new standards are the throughways by which the favored few will reach farthest into new territory—or are they?

Perhaps new standards are not essential: The Christmas 1998 shopping season proved that the central question for electronic com-

merce (E-commerce) had shifted from “whether” to “how much, how soon” (and the 1999 season was more than twice as busy). But proponents of more advanced services, such as shopping (ro)bots, effortless E-currency, or search engines with more intelligence would argue that the “netizen” ten years hence will not be able to understand how people got along in 1999 with such primitive offerings.

If the digital economy requires new standards, the process by which they are formulated and disseminated becomes central to its prospects. Will it be well-served by today’s standards processes—that is, will standards arise that are both well-conceived and timely?

This report seeks to shed some light on this question by successively discussing the place of standards (Chapter Two), lessons from five case studies (Chapter Three and Appendixes A through E), the emerging challenge of common semantics (Chapter Four), standards development institutions (Chapter Five), and public policy (Chapter Six). Chapter Seven presents conclusions, and Appendix F discusses the meaning of the term *standard*.

THE PLACE OF STANDARDS

Five years ago, Wall Street, Silicon Valley, and Hollywood hoisted competing visions of the information superhighway. Many were backed by billions of dollars, whether from bonds, venture capitalists, or ticket sales. "Set-top boxes" were a popular focus.

The Internet, by contrast, had no such backers and modest governance. But it did have standards. And that was enough to prevail.

Conceived in the 1960s, the Internet was realized in the 1970s and early 1980s with the development and refinement of protocols for message transport (Transmission Control Protocol/Internet Protocol [TCP/IP]), file transfer (File Transfer Protocol [FTP]), E-mail (Simple Mail Transfer Protocol [SMTP]), and the ability to log onto remote systems (telnet). Such standards, coupled with a spare structure for addressing (Domain Name Service [DNS]), routing, and technology insertion (the Internet Engineering Task Force [IETF]), supplied the rules by which new networks could link themselves to the Internet and thereby exchange information with users on old networks and with each other.

It took standards from outside the IETF, however, to propel the Internet into today's prominence. The development, circa 1990, of the Hypertext Markup Language (HTML) and Hypertext Transfer Protocol (HTTP) provided a foundation for creating and transferring structurally complex documents across the Internet. Once graphical browsers appeared in 1992–1993 to take advantage of these standards, the Internet became visually exciting. The existence of display tools elicited content; with content came the demand for Internet membership and yet more tools.

The Internet and the World Wide Web, as it brought together disparate threads of information technology, also affected standards. Those compatible with the Web—such as Adobe's Portable Document Format (PDF), CompuServe's Graphics Interchange Format (GIF) for images, Motion Pictures Expert Group (MPEG) music compression, and Pretty Good Privacy (PGP) encryption—did well. Those left behind by the Web—such as computer graphics metafile, the American National Standards Institute's (ANSI's) X12 for electronic data interchange (EDI) for business, the Ada programming language, and Microsoft's rich text format—did not.

To get from the *present* to the *future* relationship of standards to the digital economy, it first helps to ask what standards do.¹

WHAT MAKES A STANDARD STANDARD?

Computers, swift but stupid, are poor at inferring what something—a program, a user, another computer, a network—means, as opposed to the ones and zeroes actually used to convey data. Information and information-transfer mechanisms must therefore be composed in precise and mutually understood terms. If a convention for doing so is sufficiently common, it can be called a standard. An imprimatur of such a convention from one or another standards development organization (SDO) is not necessary but does help. Formal standards descriptions tend to be rigorous and clearly spelled out (particularly when contrasted to proprietary conventions).

A convention may be judged by its technical merits: Does it solve a problem? Does it do so elegantly? Is its solution clear? Is it easy to implement? Is it powerful enough to permit users to do what they want to do? Can its correct use be easily tested? A standard may also be judged by the fairness of the process in which it was developed: SDOs are also pickier about due process, which makes their products formally reviewed and, some believe, more fair.

¹Considerable work on standards theory was undertaken in the late 1980s and early 1990s. See David and Greenstein (1991), Spring (1991), and *Information Infrastructure and Policy's* special issue on interoperability (1995). For a broader perspective on standards and the digital economy, see Shapiro and Varian (1999).

Nevertheless, the true test of a standard is that it be widely used. The wider the use, the lower the cost of interoperation between two random users, the more people and processes that can interact with each other, the less the need for translation (and the inevitable loss of meaning) to exchange information, and the greater economies of scale in producing support services, tools, and training. Once a convention becomes a true standard, alternatives tend to lose support—leaving some users worse off (e.g., even in a world where C++ dominates, other computer languages, such as Ada, have their unique strengths). Furthermore, the best conventions do not necessarily graduate to standards: Those that win early acceptance or are merely crowned by the expectation of success may attract the next wave of users who want to interoperate with as many prior users as possible (i.e., the “network effect”) or who at least do not wish to be stranded down the road. The more users, the greater the expectations of further success. And so on.

There are essentially two approaches to standardization. *Minimalists* value simplicity and rapid uptake by the user community. Their standards tend to be expressed as primitives from which subsequent elaboration takes place *after* acceptance occurs. Theirs is an inside-out world. *Structuralists* value comprehensiveness and precision in the fear that rough-and-ready standards will, at best, grow like weeds, making well-kept ontological gardens that much harder to maintain. They would model the world so comprehensively that no human activity, extant or imagined, would fall outside their construct.² Such activities are then mapped into successively finer categories of relationships, which are then enumerated and labeled. Theirs is an outside-in world.

Open Systems Interconnection (OSI) is clearly structuralist: It grew from a reference model that partitioned all data communications into seven layers, from the physical exchange of bits to the organization of data via applications (e.g., E-mail). The OSI reference model is universally acknowledged and rarely followed as such. The Common Object Request Broker Architecture (CORBA) is another similarly ambitious (albeit less structuralist) set of standards for

²A typical structuralist approach is the Universal Modeling Language, a spin-off from the Ada computer language community.

object-oriented middleware that would rope in an enterprise's legacy base of software and hardware.

The Internet, by contrast, was created by minimalists who forswore grand conceptions to focus on a few good protocols (e.g., TCP/IP) that would permit the services they wanted. It handily beat OSI at its own game. Although the Web's creators may have sought a comprehensive structure to the universe of documents (see Berners-Lee, Connolly, and Swick, 1999), HTML rose to prominence as a set of well-chosen primitives rather than the expression of any such structure.

Successful standards, correspondingly, tend to start small, not large. The entire C reference manual fills no more than 40 pages of broadly spaced print. (Kernighan and Ritchie, 1978, pp. 179–219.) The first version of HTML could be learned, in its entirety, in an hour. TCP/IP, and the Structured Query Language's (SQL's) rules could be stated very succinctly. By contrast, very complex standards—such as OSI, Ada, Integrated Services Digital Network (ISDN), the Standard for the Exchange of Product model data (STEP)—were born large and complex; the first two have largely failed, and the second two are struggling. Of course, neither C (which grew to C++) nor HTML *stayed* simple, but both caught on before they evolved toward greater complexity.

Standards are shaped by conflicts within and between communities of computer engineers and corporate representatives. Engineers prefer standards be elegant and functional; aesthetic differences often lead to fierce standards fights arcane to outsiders. Corporations cooperate to use standards for new markets, but vie over their details to widen or narrow access to existing ones (depending on who is on top). Engineers bickered over the SONET fiber-optic trunk-line standard until corporate executives commanded them to get on with finding a standard so that firms could interconnect. Standards also permit many technical features of a product to be described in shorthand, leaving companies to play up its unique features. Engineers alone battled over whether Open Step or Motif would become the standard graphical user interface for X-Window/UNIX systems—until corporations realized that user interfaces were a useful way to differentiate workstations. Then distinctions were emphasized in public.

Browsers appear to play a key role in validating Web standards—in the sense that an innovation not supported by a browser is in trouble. For the nonce, browsers remain the door to the Web and hence the digital economy.³ Not for nothing has Microsoft's assault on the *browser* market been front and center in its antitrust case—even though this one product contributes but a small fraction of its business and makes little money on its own. But how much influence do browser companies wield in the overall process? Java caught the imagination of developers before it showed up in browsers. Yet, if Java had not shown up in one soon enough, it would have died. Conversely, once Java had enough momentum, any browser that did not support it would have hurt itself. Clearly any *new* version of Java or HTML (or its putative successor, the Extensible Markup Language [XML]) *not* supported by a popular browser has a poor chance of success. Assisting the browser is a vast array of plug-ins with conversion, display, and manipulation capabilities (some of which, such as Adobe's Acrobat software for reading PDF files, antedate the Web).⁴

THE POTENTIAL IMPORTANCE OF STANDARDS

Standards failures leave up to three problems in their wake: (1) New activities are stillborn; (2) new activities emerge but with little interoperability among domains that follow their unique standards; or (3) proprietary standards enable a thriving but biased marketplace, thereby reducing competition and ultimately retarding innovation.

Good standards clarify investment decisions. Since everyone uses TCP/IP for packaging Internet content, engineers understand what they have to engineer their networks to do. Those who generate content or support services, in turn, know what they have to break down their information streams into. Network providers need not worry so much about what kind of content they are carrying, and

³But not the only door. Popular techniques for real-time audio and video streaming do not work through browsers; neither do downloads to palmtops. Instant messaging has also been viewed as a new portal into the Web; see Paul Hagan of Forrester Research, as quoted in Ricciuti (1999).

⁴A plug-in is a piece of software that a browser loads to perform a specific function. Netscape's Web site (<http://www.netscape.com/plugin/index.html>) listed 176 external plug-ins as of March 16, 1999.

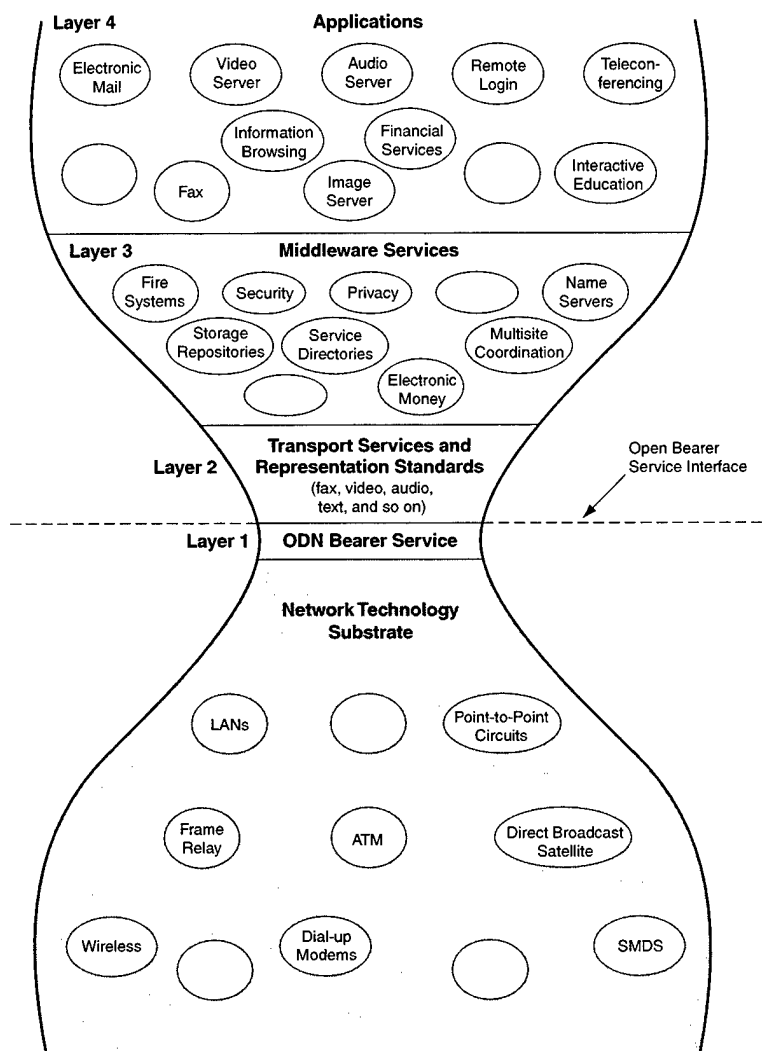
content providers do not need to worry about what networks their content flows over. Consider Figure 1. Multiple information applications, services, and formats are listed on the top, and multiple network technologies are listed on the bottom, but only one transport-cum-addressing service occupies the middle. The latter is clearly TCP/IP (and/or descendants). TCP/IP's achievement was to simplify an otherwise exceedingly complex three-dimensional mix-and-match problem into a tractable, two-dimensional mix-and-match problem.

TCP/IP also illustrates how architecture emerges from standards. Both the packet switching of TCP/IP and the circuit switching of telephony can route messages, but they lead to different kinds of networks. Circuit switching, with its limited and controlled data-stream handoffs and its parceling of bandwidth in discrete units (e.g., of 64,000 bps lines), facilitates per-use billing and system management but frustrates the carriage of bursty data flows and high-bandwidth multimedia (which require bundling and synchronizing multiple lines). The multiple and globally unpredictable handoffs of discrete TCP/IP packets complicate per-use billing and system management, but TCP/IP is tailor-made for higher bandwidth. Telephony concentrates intelligence at the switch; packet switching concentrates intelligence at the terminal. (See Isenberg, 1997.) More generally, packetization obviates worry about which bit of content (e.g., voice, video, and data) uses which internal channel (e.g., which time-slice, or *n*th bit of 16). Going farther, markup languages (such as XML) permit structured content to be expressed without worry over what position or how many bits a particular datum occupies.

Cellular telephony illustrates the power of standards even over technology. In 1983, AT&T's roll-out of an analog standard (Advanced Mobile Phone System [AMPS]) kick-started cellular telephony in the United States. But Europe, with its multiple cellular standards, reaped confusion. So vexed, and anticipating a second generation of cellular systems based on digital technology, European countries agreed to develop a common system (Glenn et al., 1999) and, in 1991, deployed Groupe Spéciale Mobile (GSM), a time-division multiple-access (TDMA) standard. Meanwhile, in the United States, the absence of a mandated TDMA standard for cellular phones allowed Qualcomm, a start-up, to introduce, in 1990, another convention for

The Open Data Network

RAND/MR1215-1



SOURCE: Reprinted with permission from National Research Council (1994), p. 53. Copyright 1994 by the National Academy of Sciences. Courtesy of the National Academy Press, Washington, D.C.

Figure 1—A Four-Layer Model for the Open Data Network

digital telephony, code-division multiple access (CDMA). It offered greater security and capacity (partially by using statistical multiplexing to exploit the fact that 60 percent of all voice circuits are silent at any one time). TDMA vendors responded with new frequency allocation methods that promised great increases in capacity.⁵ Whose philosophy won? GSM was a great fillip to cellular telephony in Europe, permitting a level of continent-wide roaming long unavailable in the United States (where older analog systems remained in use). Furthermore, because GSM is a *global* standard, whereas Qualcomm's CDMA was but one of many *national* standards, GSM phone users could roam overseas as well. As of 1998, GSM had claimed 64 percent of the world market—well over 90 percent outside North America. Cellular telephony stands out as a high-technology market unique because its major players, Ericsson and Nokia, are European. Europe also appears ahead in putting Web access on cell phones. The International Telecommunications Union (ITU) is now working on a third generation of mobile systems capable of raising bandwidth up to two megabits per second. As 1999 ended, a compromise between Qualcomm and Ericsson left a CDMA proposal as the likely choice—even though the last standard was TDMA. There is little evidence that Europe's ability to achieve market dominance based on TDMA is any bar to their being able to do so again with CDMA.

⁵See, for instance, Therrien (1992).

LESSONS FROM FIVE CASE STUDIES

Strong examples, such as the two just discussed, come from the world of communications infrastructures, where horizontal interconnection is a sine qua non of the business, and poor bets can cost companies billions of dollars. Physical infrastructure does not seem to be an E-commerce barrier, but the same may not necessarily hold for the semantic infrastructure (the encapsulating of business concepts into terms recognized by computers).

Approaching the issue of Web standards required doing four case studies on the present and one more on the future. The first examines two key components of the Web page, HTML and Java. The second focuses on XML and how markup may be used to bring order not only to the Web but also to E-commerce. The third discusses the raft of standards proposed to organize knowledge. The fourth deals with payments, privacy, and the protection of intellectual property. The last looks at the future of standards as a function of the still-evolving value chains of E-commerce.

These case studies both reinforce what decades of prior standards have already proven and acknowledge new requirements for external interoperability and a reasonable intellectual property regime. Five lessons merit attention.

STANDARDS FOSTER OPENNESS

The story of Web, E-commerce, and knowledge organization standards proves again that standards, regardless of how earnestly people try to manipulate them, are a force for openness. Where standards are absent, or ill-suited for their task, markets are closed or

constricted, and raw market power prevails. Complaints about bias in the standards process are essentially secondary unless bias prevents the standards process from functioning at all. True, many standards battles (Netscape versus Microsoft on HTML, Sun versus Microsoft on Java, America Online [AOL] versus Microsoft on instant messaging) stem from disputes between Microsoft and an opposing coalition. And neither monopoly control nor hostile bifurcation is necessarily desirable. Nevertheless, it is difficult to find a *de jure* standard, howsoever skewed in development, that enshrined a market leader as well as closed or *de facto* standards have done.

BUT STANDARDS HAVE TO SOLVE PROBLEMS, BOTH TECHNICAL AND SOCIAL, TO SUCCEED

The life span of standards can often be predicted by gauging what and whose problems they solve. Thus, Secure Electronic Transactions (SET—a payment mechanism) has lagged because consumers have not been convinced they needed its authentication services; micropayments have lagged because consumers have not been persuaded to pay for information or, at any rate, not in dribs and drabs. HTML permitted users to look at documents as they access them—something FTP alone did not provide. Java was a standard in search of a market, and once the expression of animated GIFs was standardized, it had less to offer the Web. This rule will doubtlessly apply to standards for software agents: Is this something customers need? Standards for knowledge organization illustrate the *whose* aspect: Can the needs of librarians persuade authors to categorize their works? Is the interlibrary loan model relevant to digital material? Can lawyers persuade publishers to identify the property rights inherent in a work?

THE INTERNET AND WORLD WIDE WEB HAVE SHIFTED THE FOCUS OF INTEROPERABILITY

In the 1980s and early 1990s, many firms that automated their departments separately found themselves with a large headache when building an enterprise system from them. The Standard Generalized Markup Language (SGML—a way to mark up compound documents) and CORBA (a way to build applications from components held by a network) were touted as middleware glue. Today the

emphasis is shifting to linking with external customers and suppliers. Middleware has proven too heavy for external systems; ANSI X12 is suitable only for repeat business along well-established lines. Hence the popularity of XML, which lightened SGML and does not assume the existence of middleware or even that external users will employ common practices and models. XML has started to replace CORBA as a syntactic layer for standards ranging from CommerceNet's ecoSystem, Hospital Layer 7 (HL7), and three standards from Case Study 3 (Simple Digital Library Interoperability Protocol [SDLIP], the Dublin Core, and PubMed).

LIGHT STANDARDS CONTINUE TO DO BETTER

The simplicity of HTML, Javascript, and the Secure Socket Layer (SSL) has prompted their uptake on the Web. XML, by simplifying SGML, has given markup a great lift. The Dublin Core looks light enough to succeed. By contrast, SET and many of the proposed knowledge organization standards appear too heavy for takeoff, and the complex structural models being built for RDF (resource description framework) or used to bulwark future object identifier models do not feed optimism about either.

BUT THE ENCAPSULATION OF THE REAL WORLD INTO STANDARD SEMANTICS IS LIKELY TO BE DIFFICULT

With the enthusiastic adoption of the metalanguage, XML, issues of syntax, the easy work of standardization, appear settled. The gradient ahead to semantic standards is far steeper, with no obvious trail upward. This is because semantic standards are an abstraction of a complex universe. Backers of EDI/X12, HL7, and perhaps the Dublin Core must hope the semantic structures and implicit business models of earlier standards may be converted into straight semantics. Otherwise, common *notations* overlaid upon dissimilar *notions* of how the world of discourse is constructed will lead to ambiguity: messy for humans and dangerous for machines.

Indeed, the search for semantic standards is becoming the touchstone for all upper-level standards efforts. How to reach that goal merits consideration in its own right.

THE EMERGING CHALLENGE OF COMMON SEMANTICS

With XML has come a proliferation of consortia from every industry imaginable to populate structured material with standard terms (see Appendix B). By one estimate, a new industry consortium is founded every week, perhaps one in four of which can collect serious membership dues. Rising in concert are intermediary groups to provide a consistent dictionary in cyberspace, in which each consortium's words are registered and catalogued.

Having come so far with a syntactic standard, XML, will E-commerce and knowledge organization stall out in semantic confusion? With at least one human taking part in every transaction, business-to-consumer commerce should not be greatly affected (poor prospects for shopping bots may not bother site owners that profit from strong brand loyalty). But standardization matters greatly for business-to-business commerce, with its repeat purchases, steady cost pressures, and potential savings from tying purchasing to automated production and scheduling systems. This also holds for knowledge organization, with many subject areas supported by literally millions of documents.

How are semantic standards to come about? Five paths are suggested below.

LET THE MARKET DECIDE

At first, multiple standards consortia create competing vocabularies, some better than others. Confusion reigns. Many small clusters latch onto one or another standard; others follow the dictates of their primary client. Everyone else, paralyzed by the many choices,

watches and waits. In time (perhaps only Internet time), momentum develops for a preferred tag set. Once this momentum is recognized, competing alternatives are discarded, and consolidation proceeds rapidly. Everyone ends up speaking the same language.

That is the happy version. It appropriates the advantages of natural selection in that the fittest survive, and no big brother, whether public or private, need intervene with a heavy and perhaps clumsy hand (alternatively, such intervention starts, but the fingers work so slowly that consolidation takes place before the grip is tightened).

But is the happy version likely? A chicken-and-egg cycle may yield paltry results: E-commerce remains a manual undertaking without universal standards with which to program computers, and the forces that would foster consolidated standards work without great urgency because the applications that need such standards are not imminent. E-commerce clusters may even form around a dominant buyer or vendor (e.g., for office supplies—although the Open Buying Initiative is headed by an Office Depot vice president), duplicating in cyberspace the kind of *keiretsu* that the Japanese invented for real space. Competing clusters that form at the national (or linguistic) level may, ironically, retard today's healthy progress toward a global economy. It is unclear whether such clusters would be precursors or barriers to eventual consolidation.

Granted, standards could consolidate too fast without adequate consideration of alternatives. But a greater threat arises because the ecology of standards is anything but natural. Absent standards, profits await institutions that can shepherd the bulk of transactions under their roofs for a small fee. Even a standard born of proprietary instincts may foster a monopoly over critical aspects of E-commerce. Worse may result if the winner is already a full or near monopolist in an ancillary field (e.g., office software, on-line Internet provision), so that one monopoly position reinforces another.

HAVE DIFFERENT COMMUNITIES EACH DECIDE

A variant of the Darwinian struggle is that each sector generates a common vocabulary for its own business based on its own standards work. Smaller, more homogenous groups may succeed where larger ones fail.

Would sector-specific standards suffice? The boundary between sectors has never been easy to delineate and is not always meaningful. Axles and tires are sold to automakers; tires and uniforms, to department stores; uniforms and medical services, to hospitals. Many large customers—not least, the federal government—would have to conduct E-commerce across a wide span of sectors. Because every standard rests on its own business model, which is reified in complex enterprise management software, multiple business models make business process systems unwieldy to create, operate, and maintain. The digital economy is redefining communities anyway. Who would have thought that the orderly business of bookselling and the chaotic business of auctioneering would have Internet business models with such common features? If the spirit of the Internet is universality, why settle for standards with the opposite effect?

ASSUME INTELLIGENT SOFTWARE WILL MEDIATE AMONG VARIOUS VOCABULARIES

Conceding diverse tag sets, another approach to E-commerce and knowledge organization would have sophisticated software mediate among them, much as people who speak different languages can be understood through translation.

If the standards problem were no more than a simple one-for-one substitution (“you say tomato . . .”) this approach could work. But translation presumes a common cognitive model of the universe described by various words. Uniformities on the structuring of text make it possible (if not easy) to translate between documents produced by Word Perfect and by Microsoft Word. Greater variations in the structure of graphical files make similar translation between Harvard Graphics and Microsoft PowerPoint nearly impossible. A resident of Calgary may have an easier time referring to winter in a conversation with a Quebecer, despite the difference in language, than in doing so with a Houstonian whose climate model differs greatly.

As a point of departure, the current U.S. EDI standards (ANSI X12) speak both to common business processes (e.g., invoices) and industry-specific ones (e.g., for perishables, automobile parts, and

hospital services).¹ Some business processes (e.g., invoices, again) are well-established but not in detail and not all. The simple concept of “offering price” may represent a constant in one model, a variable as a function of quantity in another, a variable as a function of time-liness in a third, and so on. An expiration date may have different meanings in food and photographic film (when discounting starts), in pharmaceuticals (when sales must end), and in software (when the sample ceases to work unless a key is purchased).

It is also unclear how tolerant business people will be for imperfect translation. Translation software is bound to be extremely sophisticated, and debugging it thoroughly may take years—and even then may not be entirely trusted.²

DEVELOP STANDARD ONTOLOGIES INTO WHICH STANDARD TERMS ARE MAPPED

Can disparate vocabularies be resolved through an ontological framework upon which each one would rest and to which each would refer? Such work is going on now (e.g., Ontology.org). Yet, finding middle ground between too little work on the area (indicating little interest) or too much work (indicating irreconcilable products at the end) is hard. Further, will the practical types that now go to standards groups be of a mind to profit from the work of the academic types that used to go to standards meetings and are still attracted by the high cognitive efforts entailed in building ontologies?

The quest for a philosophically clean language dates back before Ludwig Wittgenstein mooted the possibility in his first masterwork, *Tractatus* and conceded defeat in his last, *Philosophical Investigations*. The Defense Advanced Research Projects Agency (DARPA) has been investing in ontological development through most of the 1990s. Researchers initially optimistic about translation came to

¹But traditional EDI (as Appendix B notes) is expensive to set up; is costly to operate (especially if it requires joining a proprietary value-added network); and, as a result, was hardly universal, even at its peak.

²Tim Berners-Lee has argued that common semantics may be inferred, in part, through analytic engines that can comb the Web and see how terms are used. (See Berners-Lee, 1999, pp. 177–196.)

believe that translation was likely to be adequate only within specific domains (e.g., answering weather and travel-related inquiries)—which seem to grow narrower with every reconsideration.

CONCENTRATE ON THE KEY WORDS

A compromise is to concede that, at best, some semantic primitives will be widely understood; others will be understood within specific communities; and the rest will have to be negotiated based on commonly accessible references.

Here, the broad standards community would seek consensus on what these primitives should be and how they should be defined. If and as standards take hold, they can be expanded outward. The prospects of success may be gauged by the record of successful standards that started small and grew rather than those that were born complex.

But prospects are not guarantees. There still needs to be some forum through which agreement can be sought on *two* levels: what is to be standardized and how. It is also unproven that there is a core set of E-commerce words that is small enough to be tractable for standardization purposes, common enough among the variegated world of business models, and yet large enough to encompass most of what a minimally useful E-commerce transaction must contain.

CODA

One possible approach, which is to have the federal government drive a solution through dictate or buying power, is simply on no one's agenda. No one is asking for it, least of all those most active in the various standards processes; the government's track record of championing specific standards is, at best, uneven (e.g., continuing to back OSI as the world turned to TCP/IP); and, although the government is toward the front of the E-commerce parade, it is not *at* the front. Explorations conducted by networks of interested people scattered throughout the bureaucracy are a far cry from having a coherent policy and direction. Even European governments, historically more eager to take charge (and whose purchases account for a larger share of their region's gross domestic product) have been

holding back, waiting for their private sector to work its way toward standard.

Yet, as a practical matter, the ability of industry to develop coherent semantic standards depends on the health of the standards process.

STANDARDS DEVELOPMENT INSTITUTIONS

Entrepreneurs propose, and the standards bodies dispose, tidying up the chaotic effusions of this or that brainstorm so that mutual comprehension may reign—in theory. But is the ecology of standardization as healthy as it should be? Can standardizers still rise above the tumult of competition without ascending the sterile heights of irrelevant perfectionism? How well, in fact, *do* today's standards organizations—from United Nations (UN)-sponsored groups to ad hoc consortia—work?

The victory of the Internet over OSI in the early 1990s did lend a retrospective aura to the Internet's build-a-little, test-a-little standards processes compared to the International Organization for Standards' (ISO's) more formal habits. Some of the comparison may have been unfair: HTML and XML were based on SGML, a bona fide ISO standard, and Javascript found a home at the European Computer Manufacturers Association (ECMA), an SDO. But the common wisdom persists. Does it still hold?

THE IETF

When the tide shifted from OSI to the Internet, the attention of business shifted as well. Five to ten years ago, IETF standards meetings were dominated by academics and other computer scientists; these days, businesspeople are likely to make up the overwhelming majority of participants—even where the subject is libraries. In 1987, the IETF's semiannual meetings had only 100 attendees (up from 15 a year earlier) and for at least five years afterward hosted a community whose members knew each other. With little money at stake, partic-

ipants largely represented themselves. Agreements, while never easy, benefited from rapid feedback and ready bench-level testing of concepts.

Today, most IETF participants represent large concerns (some with stratospheric market values). Technical details are no longer so technical. Two thousand people attend the semiannual meetings. The IETF itself has been moved under the aegis of the Internet Society, which has self-consciously made itself international in recognition of often-different perspectives overseas.

Predictably, the IETF slowed down. The growing crew of network designers, operators, vendors, and researchers collectively created a bottleneck, preventing the rapid movement of standards. In theory, IETF standards processes are expeditious: six plus months for the Internet community to comment on proposed standards before they become draft standards and four plus months more until actual promotion to a standard. The effective time span is now longer. Between 1993 and 1999, it took roughly 3 years for a proposed standard to become a draft standard, and 5 years for a proposed standard to become a standard.

As Table 1 indicates, the number of proposed standards has increased threefold every three years since the mid-1980s, and number of draft standards rose similarly until the mid-1990s before level-

Table 1
Internet Proposed, Draft, and Final Standards
(by year)

	Standard	Draft Standard	Proposed Standard
-1980	1		8
1981-1983	12		2
1984-1986	11		2
1987-1989	16	4	10
1990-1992	7	10	34
1993-1995	9	40	107
1996-1998	6	31	263

NOTE: For the *latest* version of the standard.

ing off. Meanwhile, the actual total of standards is, if anything, off a bit from its 1980s pace. The IETF created a total of 40 standards between 1981 and 1989. Yet, only 22, roughly half as many, were passed between 1990 and 1998—despite the pace of technological breakthroughs and the birth of the Web.

The IETF currently has more than 100 working groups and is continuously forming new ones but with no corresponding increase in the number of standards. Despite intense debate for or against the various protocols¹ and IETF's motto of "rough consensus, running code," the requirement for rough consensus has led to splinter groups and yet more delay.

Such slow responses have sent many participants looking for a better way at a time when quick decisions are needed to keep pace with the burgeoning field of E-commerce. In 1996, a working group on Simple Network Management Protocol, Version 2 (SNMPv2) disbanded after a heated disagreement involving criteria for security and administrative standards. The ripples, felt throughout the computer-programming community, kept companies from implementing new versions of SNMPv2. Three years later, private vendors could only hope that the IETF's reassembled group will be able to endorse the 1998 SNMPv3 proposed standard. (Duffy, 1998.)

The IETF has responded by emphasizing the openness of the Internet community, the ability of that community to comment freely on issues that directly affect it, and the role of debates in weeding out inferior technology and providing technically superior standards. Recently, the debate process has seen hints of governance: As the Internet has grown, people unacculturated by their predecessors have continued to put forward their opinions in mailing lists, even after decisively hostile review; starting in 1998, people have been dropped from such lists.

Overall, the IETF has evolved away from being the progenitor of standards to the body that brings concepts into consensus. HTML and HTTP, as noted, arose from *outside* the IETF.

¹Phillip Gross, a former chairman of IETF, interview in MacAskill (1988).

ISO, ITU, AND ECMA

International standards organizations, having been overtaken by the IETF, are trying to hasten their standards processes. The ISO has adopted a Publicly Available Specification process (see Appendix A) through which standards blessed in another forum can be speeded through to final ISO imprimatur. It has yet to see wide usage.

After 1992, the ITU ended its rule that standards be approved only in Olympic years. In special cases, standards can travel from proposal to imprimatur in five months, and nine months for others is not unheard of—difficult to imagine unless the standard is fairly well cooked before it enters ITU's kitchen. Perhaps the most interesting battle shaping up concerns Internet telephony (indeed, the very name bespeaks the clash of two cultures). The IETF's Simple Internet Protocol Plus (SIPP) draft standard and the ITU's H.323 specification, while using a similar architectural model (the Internet's Real-Time Transport Protocol [RTP]), are quite dissimilar in their details.

In Europe, ECMA has evolved into a forum in which competitors to Microsoft can try to coronate a standard in ways they could not at home (see Appendix A). The Committee for European Normalization (CEN), another European SDO, is easing out of the standards business for E-commerce and is testing a new role: convening workshops to identify areas of informal agreement and best practices. Under CEN's umbrella, the European Commission has launched a project to promote a project called Electronic Commerce Open Marketplace for Industry, with workshops under way or in preparation in such areas as sanitary wares, hospital procurement, construction, and textiles (European Union [EU], 1999, p. 9). CEN will help operate the workshops and provide them neutral technological expertise.

THE WORLD WIDE WEB CONSORTIUM (W3C)

The W3C is the closest analog to the IETF in the realm of Web (as opposed to Net) standards and the dominant force in XML standards development. Founded in 1994 by Tim Berners-Lee, it has several hundred members (mostly corporations), who have to pay dues. However, as with the Internet, the primary influence is exercised within the various working groups, which create and publish technical specifications. When officially approved by the W3C process, these specifications are considered tantamount to official standards.

The W3C has taken at least one media hit because attendance at its annual meetings has flagged. (Garfinkel, 1998.) However, it is firmly in control of its realm—the development of syntactic standards. Such standards ride atop the more bit-oriented standards of the IETF, although the exact boundary between the two is undefined (e.g., who builds the next version of HTTP). Meanwhile, the job of building semantic standards to exploit the W3C's syntactic standards is the province of consortia, such as OASIS.

THE WIRELESS ACCESS PROTOCOL (WAP) FORUM

The ecology of standards is populated by start-up consortia of multiple sizes and various life spans. The 147-member² WAP Forum illustrates some typical features. It was founded in 1997 to foster the use of browsers for cell phones. The world's big-three cell-phone makers—Ericsson, Nokia, and Motorola—were founders; a fourth, Phone.com, is a start-up that actually wrote the standards. The forum's literature emphasizes that it is not a standards group (but it has a three-stage specification approval process), that it will in due course submit its recommendations to SDOs, and that it liaises with SDOs and non-SDOs (e.g., the IETF, the W3C) alike. U.S. companies constitute less than half of the forum, and its standards style reflects this. Standards are layered (as were OSI's), are middleweight in complexity, and reflect key architectural assumptions (e.g., that cell phones have keypads, receive information in "cards," but do not talk directly to Web servers). Palm Computing, whose Palm VII has a different architectural model, was a notable latecomer to the group.

OPEN-SOURCE SOFTWARE

The last few years has seen the rise of open-source software, notably Linux (a UNIX-like operating system with roughly a third of all Web server operating systems), Apache (Web server software that has just over half of its market), and Mozilla (the open incarnation of Netscape's browser). If popular, their presence may complicate the process by which dominant firms leverage monopoly control of key software to further close off upstream applications from competition.

²As of November 23, 1999; see the forum's Web site (<http://www.wapforum.org>).

Is open sourcing a standards process per se? The source code of Linux embodies the language as a standard.³ Its kernel has many contributors whose proposals are rigorously and enthusiastically vetted by peers from around the world. In the end, however, one person (Linus Torvalds) decides what is included. Outside the kernel, the process is more diffuse. Usually, the person who intuits the need for this or that extension gets to decide its contents (but, again, in open forum). Social mechanisms limit the degree of forking (two groups with incompatible approaches to a problem). Even so, fights have taken place between KDE and Gnome over which becomes the preferred user interface within the Linux community.

Open-source software is also no ironclad guarantee against market power. A friendly user interface, reliable hooks to the rest of a user's system, and hand-holding still play large roles in selling software. Red Hat Software has the largest share of the U.S. commercial market (even if many copies of Linux are downloaded for free) and enough "mindshare" to charge premium prices as well as to launch a successful initial placement offering (IPO).

Open-source software has a distinct advantage in that it permits users to modify the operating system to their specific needs. This is of special relevance to the federal government (notably the Department of Defense [DoD]) in its search for greater information security: Open sourcing not only allows bugs to be fixed quickly but permits an institution to release to its users a standard version with tempting capabilities removed and all the controls set correctly.

A TYPOLOGY

Open-source software raises a larger question: To what extent can open transparency in the creation of a standard substitute for more legal definitions of *fair* and *democratic*? Table 2 is a two-by-two typology of standards institutions. It suggests that having a strong

³Eric Raymond, admittedly a partisan, has compiled three papers on Linux: "The Cathedral and the Bazaar" (Raymond, 1998a), "Homesteading the Noosphere" (Raymond, 1998b), and "Open Source Software: the Halloween Document" (Raymond, 1998c).

Table 2
A Typology of Standards Organizations

	Democratic	Strong Leader
By membership	ISO	W3C
Open to all	IETF	Linux open source

leader may be an advantage, but so, to a lesser extent, is a process open to all (although the W3C is a membership organization, it does publish its draft specifications for open comment).

Standards development is a very pragmatic process, and there appears to be little barrier to new forms arising as need dictates. Ten years ago, virtually everything of note was done through formally established standards development organizations. Three new forms now vie for contention—plus the evanescent small-group consortia growing prominent in the E-commerce arena.

It would seem that the production of standards is the metric by which standards groups should be measured. But in times past, standards groups, notably those of a structuralist bent, also built intellectual foundations for subsequent standards, often generated by others. The OSI reference model is an example. Ontologies may provide a future one. Their value may be a question of timing. Information technology, like many human endeavors, undergoes cycles of efflorescence (from the early 1940s to the early 1960s and from the early 1980s onward) followed by consolidation (from the early 1960s to the early 1980s). A new wave of consolidation is inevitable. When its time comes, the existence of reference models may help promote standardization. But if standards development organizations do not do the intellectual spadework, who will? Support from DARPA or the National Institute of Standards and Technology (NIST) (see below) to develop a common architectural vision for future standards development may, at some point, be called for.

A healthy standards process helps foster good standards but, in the end, cannot guarantee them. Standards that prevail in the market necessarily reflect market forces: the desires of consumers; the strategies of players; and, yes, the role of public policy.

THE PLACE OF STANDARDS

The administration's objectives for E-commerce include expanded markets, protected privacy, antitrust enforcement, fraud prevention, and the protection of intellectual property rights. Clinton and Gore (1997) cited standards as critical because

they can allow products and services from different vendors to work together. They also encourage competition and reduce uncertainty in the global marketplace. Premature standardization, however, can "lock in" outdated technology. Standards also can be employed as de facto non-tariff trade barriers, to "lock out" non-indigenous businesses from a particular national market.

In theory, standards are a phenomenon of business that the government could easily stay well away from. They are very technical details of a technical enterprise. Thirty to forty years ago, when the information technology arena was smaller and less central to the overall economy, the federal role in and influence over the field was larger. Many federal users were leading-edge consumers. Then, an activist federal policy in developing standards made sense. Today, the government is but one user and not notably ahead of others.

Indeed, are standards a problem that demands government intervention? Arguing that a lack of standards is depressing *today's* growth curves is not easy. Business-to-consumer curves are growing nicely,¹ and market valuations of Internet companies are unprecedented. One might as well argue that "Internet time" is too slow.

¹The growth curve of business-to-business E-commerce is less dramatic but still impressive. After all, such E-commerce has been under way for close to 20 years. There is thus a working model and a set of relationships that need only be electrified.

But cultivating a blissful ignorance of how standards are evolving and what such evolution may mean to the market is unwarranted. First, government policymaking will affect the standards process in any case. It helps to understand how standards work if deleterious effects of such policies are to be minimized. Second, standards have the potential to become a policy tool in their own right (and, if advocated as such by the United States in international forums, must be done in a well-informed manner). Third, there is an important role to be played by a neutral, third-party convener for standards.

THE PATENT TRAP

Generally speaking, the government has smiled upon standards groups, going so far as to devote a large share of NIST to fostering their success. Although antitrust objections might theoretically have been raised to thwart their work,² no such objections ensued in practice. The deliberations of the aforementioned focus group, however, indicated that no other obstacle so vexes the standards world as the growing specter of software patents.³

One focus group member argued that the

patent situation is approaching the edge of insanity in terms of what can be patented and incompetence in terms of the lack of awareness about prior art and obviousness; claimants are coming out of the woodwork with patents perceived as trivial by many computer scientists . . . [they] are becoming a major threat to our ability to standardize, develop infrastructure for, or even advance information technology.

Another observed that the patent situation is intolerable and that standards groups are unable to do their work. He added that whatever difficulty being caused by actual patents has paled before the "fear, uncertainty, and doubt that has been interjected into the game," further noting that the prospects for litigation have created an incentive to write unclear patents that disguise the potential for

²Indeed, the Open Software Foundation, a UNIX standards group, was sued by Addamax on the theory that its deliberations reduced the market prospects for the plaintiff. But Addamax lost.

³Tim Berners-Lee, a focus group member, put his objections in print in Berners-Lee (1999), pp. 196–198.

violation until it is too late. Many of these patents were not novel but extensions from the physical world into the virtual. A third found that the patent situation has

made it difficult to write standard disclosure and licensing agreements that will assure large companies that their ancillary patents will not be infringed upon [and] they are a [particular] barrier for small firms which lack both expertise and a large pool of patents of their own.

This concern was echoed by a fourth focus group member. Patents, in general, were perceived as an unnecessary burden on the process of standards formation—not so great as to stop something that everyone really needed (e.g., XML) but enough to halt progress on ancillary agreements (e.g., the Platform for Privacy Preferences Project [P3P]) that collectively fertilize the Web's rich ecology.

As one reaction, the IETF, which used to pull a proposed standard automatically if someone claimed patent rights on it, now lets working groups decide on their own. As another, the Federal Trade Commission enjoined Dell Computer from collecting license fees from users of a computer-bus standard that incorporated patentable technology, the existence of which was not revealed to standards writers until they finished. By contrast, the W3C lacks a policy on patenting technology developed jointly by its members. Some have thus used their patents to access other companies' technologies, as well as for license fees. For instance, without announcing its application, Microsoft received a patent for "cascading style sheets," which covers not only HTML but also XML. Microsoft has said that it will give away the patented technology for free, in exchange for access to patents by other companies—but has still upset its colleagues in the W3C, many of whose members have no declared intention of limiting the use of their patents only to ensure technology interchange. More notorious was a company that allegedly attended the deliberations of the P3P discussions while simultaneously preparing a patent (whose content was continually amended as discussions progressed) and springing it upon members as the P3P effort was concluding.

Patents may do more than interfere with standards development: They are two different ways of organizing markets for meeting new opportunities—and thus competition within such markets. How, for

instance, should digital content, such as music, be protected from piracy? One proposal (the Secure Digital Music Initiative [SDMI]; see Appendix D) would encode music content in a standard form, around which music player software would be written. Another would use a proprietary formula (e.g., from InterTrust, an SDMI member) protected by patents—whose success or even presence may block or blunt the ability of a public standard to evolve toward greater functionality. How can reviews of products (e.g., books) be found, given that one document cannot easily point to another written later? One method may be a proprietary service using patented business process methods. Another may be through a standard way of tagging reviews to material so that a search engine can scoop up all the reviews. Take comparison shopping—such as rating colleges. Will it be a service hosted by a patent-rich Web site or, as Appendix E suggests, an application that any (shopping ro)bots can do if college characteristics are described in standard ways? True, standards and proprietary advantage are hardly antithetical. CommerceOne, a business-to-business E-commerce firm with a high market valuation (as of the end of 1999) was spun off from CommerceNet, a standards-promoting entity. Many standards (notably in telecommunications) have incorporated patents licensed on a “fair and reasonable” basis. But the Web’s easy ubiquity may enable sites to exploit patents to obviate the need for standards (service is just a click away). Power to the patentee will give them a disproportionate top-down ability to plan the evolution of technology, even one built for an Internet that evolved from bottom-up experimentation. The result of this tension will tell how public the Web will be.

STANDARDS AS A POLICY TOOL

Standards may be able to take over much of the work otherwise required from politically controversial regulation—especially in the digital age. The evaluation of digital material can become easier and more automatic than the evaluation of written material (e.g., how many people read labels on canned goods?). Standards can be a mechanism that shifts oversight authority from government (if national governments are even the right level for a *World Wide Web*), not to overworked consumers, but into software whose parameters permit the careful consideration of categories rather than the rushed judgment over each case.

The role of voluntary standards as a substitute for regulation is epitomized by the voluntary rating system adopted in the 1960s (succeeding the Hollywood Code). It gave adults some hint about a movie's content and helped regulate what children saw. Although imperfect,⁴ it seems to work. Parallel efforts for television combine content ratings with electronic locks (V-chips) that read the rating signal—hence the need for standards—and can screen out certain shows. With the Communications Decency Act declared unconstitutional, parental regulation of Internet content relies on purchasing nanny software, which ought to grow more effective if and when rating systems for Internet content mature. It remains to be seen whether potential ratings systems for *privacy* will put the issue to bed and avoid the need for a European-like solution.

Nevertheless, if software is to evaluate content, there must be standards by which content can be automatically recognized and its veracity vouched for (a point understood by TRUSTe, as noted in Appendix D; it not only writes standards of behavior for Web sites but monitors compliance against them). Perhaps third-party reviews (e.g., how one knows that a product is “safe and effective”) can become an important part of the voluntary regulatory process. For this method to succeed, it has to be easy to find such reviews—a problem when links from review to product are straightforward, but links from product back to review are not. A standard set of back links (implemented in, e.g., XML) may permit search engines to list reviews of any well-labeled object without the usual clutter of extraneous material.

In enhancing security, notably infrastructure security, the federal government is uncomfortably seated between controversial regulation and ineffective hectoring. Standards by which institutions may have their security policies and practices rated by third parties might provide a lever for improvement. As ISO 9000 seems to have done for quality control and ISO 14000 for environmental management, the steady pressure of outside review, both good and bad, has a way of pushing people to adopt good practices, if only defensively.

⁴Witness the hybrid PG-13 and NC-17 ratings, the complaints that violence is accorded softer treatment than sex, and the irony that Hollywood feels that G-ratings dampen ticket sales.

If the government is to exploit standards as a substitute for policy, it may need to put some money into their development. A clear consensus of the focus group was that funding agencies (e.g., the National Science Foundation [NSF]) look kindly on researchers using part of their grant money to participate in the standards process. With open access to a well-structured universe of knowledge critical to its advancement, research and development on conceptual frameworks (e.g., ontologies) may also merit support.

Does standards advocacy need a standards advocate per se? The issues just noted—e.g., privacy, patents, and security—are debated at the cabinet and subcabinet levels. But the highest-ranking individual that deals with standards on a regular basis is the head of NIST's Information Technology Laboratory (ITL)—three levels below the Secretary of Commerce (via the Under Secretary for Technology Administration and the head of NIST). In national security affairs, elevating an issue entails using a chair in the National Security Council (NSC). There is no easy analog on the domestic side. Neither the National Economic Council nor the Office of Science and Technology Policy have the longevity of the NSC, and neither is involved in day-to-day management or closely linked to DoD (whose acquisition clout matters). Furthermore, raising an issue has a concreteness absent from raising a viewpoint that applies across issues, particularly one with such broad reach.

NIST'S EVOLVING ROLE

ITL has long been involved in the standards process, adding its technical expertise or good offices to promote standardization and representing the government's interest in this or that feature. To the extent that the government has sought an interoperability strategy (e.g., the DoD's Joint Technical Architecture), NIST has helped compile lists of relevant standards and profiles thereof.

After the 1994 elections, ITL expected criticism for spending tax dollars doing standards work that private firms could do and did.⁵ So, it decided to focus on the development of public infrastructure: tests,

⁵References to ITL prior to 1995 are to its predecessor, the National Computer Systems Laboratory, renamed when it absorbed the Computing and Applied Mathematics Laboratory that year.

metrics, and corpora by which adherence to standards could be measured.⁶ It was the general consensus of the focus group that even more metrology was needed. Indeed, it often happens that the technical quality of a standard cannot be truly known until efforts to write a test for it are under way.

A second area in which NIST may make a contribution is to develop (or at least evaluate) technology that may facilitate interoperability. Are there generalizable features of standards that make them easy to implement, flexible against unknown changes in requirements, and clear enough to minimize the amount of hand-massaging required to make two compliant applications interoperable? Are there ways of ensuring interoperability with lighter standards that do not have to specify as much, or, better yet, with translators and mediators that can dispense with many higher-level standards altogether? DARPA is funding some technology, and more sustained efforts may be merited. NIST itself can develop the parameters, corpora, tests, and test-beds that help measure the quality and fitness of ontologies and mediators.

A third area where NIST could play a role is in the development of a terrain map for E-commerce and knowledge-organization standards.⁷ Battered by a blizzard of standards activities—especially in the realm of semantic standards and resulting tag sets—many such enterprises have little inkling of who is doing what to whom or, ultimately, where the real action is taking place. ITL could provide a neutral meeting ground for various efforts; it could also document the current status and plans of the various groups to look for possible gaps, overlap, and contradictions. Erecting a reference structure may help standards processes and reduce overall coordination costs by letting everyone know where they sit in the rapidly evolving universe.⁸

⁶A *corpora* is a sample (e.g., ten hours of telephone conversation) against which technologies (e.g., speech recognition) are tested. NIST's testing role involves testing for *performance* as well as *conformance*.

⁷Much as NIST's Electrical and Electronics Laboratory helped Sematech.

⁸The line between a reference structure and a recommendation is thin but critical. In the mid-1990s, NIST created profiles of standards selected for having an SDO imprimatur. Nevertheless, NIST's corporate history—embodied in its defining controversy over battery-life extenders—has made it reluctant to evaluate specific products; it

FOR FURTHER RESEARCH

Both the patents issue and the uphill climb to semantic standards merit more thought.

Fundamentally, what *should* be patentable (e.g., software frameworks, business models)? How can greater awareness about prior art be introduced into the Patent and Trade Office? What standards of obviousness should be applied (e.g., are virtual metaphors for an existing physical construct *per se* patentable)? Should the patent application and enforcement process be more transparent (i.e., to reduce “fear, uncertainty, and doubt”)? How can incentives to write unclear patents be reduced? As for their effect on standards processes: Can better patent disclosure and licensing agreements be written? Can companies be asked to disclose patent applications as a precondition of joining standards efforts? Is there a way to get a quick read of how standards processes may be jeopardized by specific patent claims? How might strategies of patent makers and patent takers evolve within the standards arena?

The research on semantic standards presumes a public interest in timely standards that neither split the user community into disparate camps nor bias the marketplace in anyone’s favor. How can progress toward such an end be measured? What are the precursors of success or failure? What indicates that forking is appropriate (e.g., different standards matched to different needs) or inappropriate (e.g., differences among standards that reduce interoperability more than they improve the fit)? What determines whether standards solve the right problems, and how can this be judged by creators, consolidators, and users of content? Finally, are there techniques by which the health of standards *processes* can be judged?

validates the capability of independent laboratories to do this. It is even more loathe to handicap winners in standards contests.

Will the digital economy be well-served by standards? So far, standards difficulties have proven no worse than a speed bump before technology's relentless march. And standards developers have showed considerable flexibility in finding forums in which to generate agreement. But two issues loom as potential roadblocks.

One is the deleterious influence of patents on the standards process—a problem that standard developers may find ways to mitigate for their own purposes but that, ultimately, has to be resolved outside the standards process.

The other is the challenge of semantic standards. Here, too, the root of the problem is no less than the long-standing difficulty of encapsulating the messy world of human affairs into a clean form suitable for computing machines (or, before 1950, for mathematics). Perhaps this problem will never be solved. Or, perhaps, men and women of good will can find their way to a standard set of good-enough resolutions. Or, just maybe, someone will develop an approach simple and satisfying enough to become a standard on its own—so much so that people looking back will wonder why the problem ever existed.